

MySQL 8 Observability

(Also Query Troubleshooting with PMM Demo)

Peter Zaitsev
CEO Percona
May 17th, 2021



HighLoad++
Всё на 2021



Why Observability

**Many Non Easily
Repeatable Issues
in Complex
Systems**

**Resolving and
Preventing Issues
Requires
Understanding**

Data Capture

Ongoing Data Capture (Monitoring)

Temporary Data Capture (Debugging)

Comprehensive View Needed

You Can't just use at MySQL Alone

OS Issues, Hardware Issues are often root cause

Application issues can't be ignored

Background Load

MySQL 8 at Focus

MySQL 8 Data Sources

SHOW STATUS

**INFORMATION
SCHEMA**

**PERFORMANCE
SCHEMA**

LOGS

EXPLAIN

**OPTIMIZER
TRACE**

SHOW [GLOBAL] STATUS

Existed Forever

Shows 400+ status variables

Most are counters, some are gauges, some text

Session and Global Scope

Also available as Performance Schema Table

<https://dev.mysql.com/doc/refman/8.0/en/server-status-variables.html>

Global And Session

```
mysql> select * from performance_schema.global_status where variable_name='Questions';
```

VARIABLE_NAME	VARIABLE_VALUE
Questions	82457893

```
1 row in set (0.01 sec)
```

```
mysql> select * from performance_schema.session_status where variable_name='Questions';
```

VARIABLE_NAME	VARIABLE_VALUE
Questions	153

```
1 row in set (0.00 sec)
```


Trust but Verify

```
mysql> select * from performance_schema.session_status where  
variable_name like "%rows%";
```

VARIABLE_NAME	VARIABLE_VALUE
Innodb_rows_deleted	14553828
Innodb_rows_inserted	15840851
Innodb_rows_read	1927758552
Innodb_rows_updated	29290781
Mysqlex_rows_sent	0
Not_flushed_delayed_rows	0
Sort_rows	19

```
7 rows in set (0.00 sec)
```

Status-Profiling Query

```
mysql> select * from performance_schema.session_status where variable_name like "%Handler_read_next%";
```

```
+-----+-----+
| VARIABLE_NAME | VARIABLE_VALUE |
+-----+-----+
| Handler_read_next | 11           |
+-----+-----+

1 row in set (0.00 sec)
```

```
mysql> select count(*) from orders1;
```

```
+-----+
| count(*) |
+-----+
|    75653 |
+-----+

1 row in set (0.02 sec)
```

```
mysql> select * from performance_schema.session_status where variable_name like "%Handler_read_next%";
```

```
+-----+-----+
| VARIABLE_NAME | VARIABLE_VALUE |
+-----+-----+
| Handler_read_next | 75664         |
+-----+-----+

1 row in set (0.00 sec)
```

VMSTAT Like output for MySQL

```
root@mysql3:~# mysqladmin extended -i1 -r | grep Questions
| Questions | 82525506
| Questions | 378
| Questions | 519
| Questions | 591
```

With Percona Toolkit

```
root@mysql3:~# pt-mext -r -- mysqladmin ext -i1 -c5
```

...

Bytes_received	13701143412	54076	101332	44687
Bytes_sent	19638518220	98678	131670	94501

...

Innodb_buffer_pool_pages_data	27615	-58	-3	2
Innodb_buffer_pool_pages_dirty	13643	-58	86	-26
Innodb_buffer_pool_pages_flushed	23741158	182	145	123
Innodb_buffer_pool_pages_free	964	60	0	-2

INFORMATION_SCHEMA

60+ Tables

Some are Schema Related

Others are Performance Statistics

INFORMATION SCHEMA – Schema Info

```
mysql> select * from TABLES limit 1 \G
***** 1. row *****
TABLE_CATALOG: def
TABLE_SCHEMA: information_schema
TABLE_NAME: CHARACTER_SETS
TABLE_TYPE: SYSTEM VIEW
ENGINE: NULL
VERSION: 10
ROW_FORMAT: NULL
TABLE_ROWS: 0
AVG_ROW_LENGTH: 0
DATA_LENGTH: 0
MAX_DATA_LENGTH: 0
INDEX_LENGTH: 0
DATA_FREE: 0
AUTO_INCREMENT: NULL
CREATE_TIME: 2019-05-27 20:33:04
UPDATE_TIME: NULL
CHECK_TIME: NULL
TABLE_COLLATION: NULL
CHECKSUM: NULL
CREATE_OPTIONS:
TABLE_COMMENT:
1 row in set (0.00 sec)
```

INFORMATION SCHEMA – Metrics

```
mysql> select * from innodb_metrics limit 1 \G
***** 1. row *****
      NAME: metadata_table_handles_opened
    SUBSYSTEM: metadata
      COUNT: 20040
    MAX_COUNT: 20040
    MIN_COUNT: NULL
    AVG_COUNT: 0.10970542669608203
    COUNT_RESET: 20040
    MAX_COUNT_RESET: 20040
    MIN_COUNT_RESET: NULL
    AVG_COUNT_RESET: NULL
    TIME_ENABLED: 2019-06-11 19:01:11
    TIME_DISABLED: NULL
    TIME_ELAPSED: 182671
    TIME_RESET: NULL
      STATUS: enabled
      TYPE: counter
    COMMENT: Number of table handles opened
1 row in set (0.00 sec)
```

Note on Innodb Metrics

- **While Looks similar to SHOW STATUS is not enabled by default**
 - `innodb_monitor_enable=all`

Performance Shema

The “Instrumentation of Choice” in MySQL

100+ Tables

Some information Capture Enabled by Default

Can enable more instrumentation permanently or temporary

Overhead can be high

Performance Schema Configuration

```
mysql> show tables like "setup%";
```

```
+-----+
| Tables_in_performance_schema (setup%) |
+-----+
| setup_actors                          |
| setup_consumers                       |
| setup_instruments                     |
| setup_objects                         |
| setup_threads                         |
+-----+
5 rows in set (0.02 sec)
```

Actors – What users should be profiled ?

```
mysql> select * from setup_actors;
```

HOST	USER	ROLE	ENABLED	HISTORY
%	%	%	YES	YES

1 row in set (0.00 sec)

What Summaries should be produced

```
mysql> select * from setup_consumers;
```

+-----+-----+	
NAME	ENABLED
+-----+-----+	
events_stages_current	NO
events_stages_history	NO
events_stages_history_long	NO
events_statements_current	YES
events_statements_history	YES
events_statements_history_long	NO
events_transactions_current	YES
events_transactions_history	YES
events_transactions_history_long	NO
events_waits_current	NO
events_waits_history	NO
events_waits_history_long	NO
global_instrumentation	YES
thread_instrumentation	YES
statements_digest	YES
+-----+-----+	

```
15 rows in set (0.00 sec)
```

Instruments - Instrumentation Points

- 1200+ Instruments
- 800 enabled by default and almost 300 timed

```
mysql> select * from setup_instruments limit 2 \G
***** 1. row *****
      NAME: wait/synch/mutex/pfs/LOCK_pfs_share_list
     ENABLED: NO
      TIMED: NO
  PROPERTIES: singleton
 VOLATILITY: 1
DOCUMENTATION: Components can provide their own performance_schema tables. This lock protects the list of such tables definitions.
***** 2. row *****
      NAME: wait/synch/mutex/sql/TC_LOG_MMAP::LOCK_tc
     ENABLED: NO
      TIMED: NO
  PROPERTIES:
 VOLATILITY: 0
DOCUMENTATION: NULL
2 rows in set (0.00 sec)
```

What Objects to Instrument

```
mysql> select * from setup_objects;
+-----+-----+-----+-----+-----+
| OBJECT_TYPE | OBJECT_SCHEMA | OBJECT_NAME | ENABLED | TIMED |
+-----+-----+-----+-----+-----+
| EVENT       | mysql         | %          | NO      | NO     |
| EVENT       | performance_schema | %          | NO      | NO     |
| EVENT       | information_schema | %          | NO      | NO     |
| EVENT       | %             | %          | YES     | YES    |
| FUNCTION    | mysql         | %          | NO      | NO     |
| FUNCTION    | performance_schema | %          | NO      | NO     |
| FUNCTION    | information_schema | %          | NO      | NO     |
| FUNCTION    | %             | %          | YES     | YES    |
| PROCEDURE   | mysql         | %          | NO      | NO     |
| PROCEDURE   | performance_schema | %          | NO      | NO     |
| PROCEDURE   | information_schema | %          | NO      | NO     |
| PROCEDURE   | %             | %          | YES     | YES    |
| TABLE      | mysql         | %          | NO      | NO     |
| TABLE      | performance_schema | %          | NO      | NO     |
| TABLE      | information_schema | %          | NO      | NO     |
| TABLE      | %             | %          | YES     | YES    |
| TRIGGER     | mysql         | %          | NO      | NO     |
| TRIGGER     | performance_schema | %          | NO      | NO     |
| TRIGGER     | information_schema | %          | NO      | NO     |
| TRIGGER     | %             | %          | YES     | YES    |
+-----+-----+-----+-----+-----+
20 rows in set (0.01 sec)
```

What Threads do we want to Instrument

```
mysql> select * from setup_threads;
```

NAME	ENABLED	HISTORY	PROPERTIES	VOLATILITY	DOCUMENTATION
thread/performance_schema/setup	YES	YES	singleton	0	NULL
thread/sql/bootstrap	YES	YES	singleton	0	NULL
thread/sql/manager	YES	YES	singleton	0	NULL
thread/sql/main	YES	YES	singleton	0	NULL
thread/sql/one_connection	YES	YES	user	0	NULL
thread/innodb/srv_error_monitor_thread	YES	YES		0	NULL
thread/innodb/srv_lock_timeout_thread	YES	YES		0	NULL
...					
thread/myisam/find_all_keys	YES	YES		0	NULL
thread/mysqlx/acceptor_network	YES	YES		0	NULL
thread/mysqlx/worker	YES	YES		0	NULL
thread/sql/slave_io	YES	YES	singleton	0	NULL
thread/sql/slave_sql	YES	YES	singleton	0	NULL
thread/sql/slave_worker	YES	YES	singleton	0	NULL

```
49 rows in set (0.00 sec)
```

Example of Performance Schema Data

```
mysql> select * from events_statements_current limit 1 \G
```

```
***** 1. row *****
```

```
  THREAD_ID: 9789
```

```
  EVENT_ID: 49
```

```
END_EVENT_ID: 49
```

```
  EVENT_NAME: statement/sql/show_status
```

```
    SOURCE: init_net_server_extension.cc:95
```

```
TIMER_START: 185174708798157000
```

```
TIMER_END: 185174709926911000
```

```
TIMER_WAIT: 1128754000
```

```
  LOCK_TIME: 98000000
```

```
  SQL_TEXT: SHOW GLOBAL STATUS
```

```
    DIGEST: 070e38632eb4444e50cdcbf0b17474ba801e203add89783a24584951442a2317
```

```
  DIGEST_TEXT: SHOW GLOBAL STATUS
```

```
CURRENT_SCHEMA: NULL
```

```
...
```

```
  NESTING_EVENT_TYPE: NULL
```

```
  NESTING_EVENT_LEVEL: 0
```

```
    STATEMENT_ID: 84449971
```

```
1 row in set (0.00 sec)
```


Too Hard to Use ?

Meet Sys Schema

Sys Schema

**Views and Stored
Procedures to access
Performance Schema
Data**

Views for Humans

```
mysql> show tables;
```

```
+-----+
| Tables_in_sys |
+-----+
| host_summary |
| host_summary_by_file_io |
| host_summary_by_file_io_type |
| host_summary_by_stages |
| host_summary_by_statement_latency |
| host_summary_by_statement_type |
| innodb_buffer_stats_by_schema |
| innodb_buffer_stats_by_table |
| innodb_lock_waits |
| io_by_thread_by_latency |
| io_global_by_file_by_bytes |
```

```
...
```

Views for Machines

```
mysql> show tables like "x$%";
```

```
+-----+
| Tables_in_sys (x$%) |
+-----+
| x$host_summary      |
| x$host_summary_by_file_io |
| x$host_summary_by_file_io_type |
| x$host_summary_by_stages |
| x$host_summary_by_statement_latency |
| x$host_summary_by_statement_type |
| x$innodb_buffer_stats_by_schema |
| x$innodb_buffer_stats_by_table  |
| x$innodb_lock_waits             |
| x$io_by_thread_by_latency       |
| x$io_global_by_file_by_bytes   |
```

Example for Humans

```
mysql> select * from session limit 1 \G
***** 1. row *****
      thd_id: 7101
      conn_id: 7060
      user: app2@li1306-116.members.linode.com
      db: tpcc4
      command: Sleep
      state: NULL
      time: 7
current_statement: NULL
statement_latency: NULL
      progress: NULL
      lock_latency: 0 ps
      rows_examined: 0
      rows_sent: 0
      rows_affected: 0
      tmp_tables: 0
      tmp_disk_tables: 0
      full_scan: NO
      last_statement: COMMIT
last_statement_latency: 97.95 us
      current_memory: 119.71 KiB
      last_wait: NULL
last_wait_latency: NULL
      source: NULL
      trx_latency: 3.10 ms
      trx_state: COMMITTED
      trx_autocommit: NO
      pid: 14124
      program_name: NULL
1 row in set (0.23 sec)
```

Example for Machines

```
mysql> select * from x$session limit 1 \G
***** 1. row *****
      thd_id: 9122
      conn_id: 9081
        user: app1@li1306-116.members.linode.com
         db: tpcc5
      command: Sleep
        state: NULL
         time: 11
current_statement: NULL
statement_latency: NULL
   progress: NULL
   lock_latency: 0
  rows_examined: 0
    rows_sent: 0
  rows_affected: 0
    tmp_tables: 0
 tmp_disk_tables: 0
    full_scan: NO
   last_statement: COMMIT
last_statement_latency: 1945326000
   current_memory: 172047
    last_wait: NULL
last_wait_latency: NULL
      source: NULL
   trx_latency: 9264917000
    trx_state: COMMITTED
  trx_autocommit: NO
        pid: 17713
   program_name: NULL
1 row in set (0.55 sec)
```

Information by Client Hosts

```
mysql> select * from host_summary \G
***** 1. row *****
      host: li1306-116.members.linode.com
    statements: 58027942
  statement_latency: 9.56 h
statement_avg_latency: 592.85 us
      table_scans: 296
        file_ios: 67522180
    file_io_latency: 1.45 h
current_connections: 20
  total_connections: 148
      unique_users: 2
    current_memory: 18.39 MiB
total_memory_allocated: 2.72 TiB
```

File IO by Hosts

```
mysql> select * from host_summary_by_file_io;
```

host	ios	io_latency
background	110954519	7.63 h
li174-63.members.linode.com	33872913	43.80 m
li1306-116.members.linode.com	33795434	43.66 m
li534-125.members.linode.com	33767097	43.08 m
localhost	1138434	1.23 m

```
5 rows in set (0.01 sec)
```


Types of IO

```
mysql> select * from host_summary_by_file_io_type;
```

host	event_name	total	total_latency	max_latency
background	wait/io/file/innodb/innodb_log_file	87369393	7.22 h	430.95 ms
background	wait/io/file/innodb/innodb_data_file	22990098	16.77 m	134.75 ms
background	wait/io/file/innodb/innodb_parallel_dblwrite_file	387716	7.67 m	115.12 ms
background	wait/io/file/sql/binlog	232529	2.09 s	28.38 ms
background	wait/io/file/innodb/innodb_temp_file	415	49.44 ms	2.32 ms
background	wait/io/file/sql/slow_log	7	5.05 ms	4.58 ms
background	wait/io/file/sql/pid	3	3.02 ms	2.97 ms
background	wait/io/file/sql/ERRMSG	5	2.05 ms	1.04 ms
background	wait/io/file/sql/binlog_index	36	1.71 ms	643.66 us
background	wait/io/file/mysys/cnf	5	1.58 ms	1.27 ms
background	wait/io/file/sql/casetest	15	1.24 ms	1.04 ms
background	wait/io/file/mysys/charset	3	426.80 us	408.13 us
background	wait/io/file/sql/misc	1	48.61 us	48.61 us
li1306-116.members.linode.com	wait/io/file/sql/binlog	3074720	19.91 m	116.84 ms
li1306-116.members.linode.com	wait/io/file/innodb/innodb_data_file	1678043	15.42 m	173.48 ms
li1306-116.members.linode.com	wait/io/file/sql/slow_log	29049102	8.34 m	75.98 ms
li1306-116.members.linode.com	wait/io/file/sql/binlog_index	165	7.71 ms	2.32 ms

Statement Summary Per Host

```
mysql> select * from host_summary_by_statement_latency \G
***** 1. row
host: li174-63.members.linode.com
      total: 29175532
total_latency: 4.80 h
  max_latency: 51.14 s
  lock_latency: 37.47 m
    rows_sent: 13377371
rows_examined: 686152549
rows_affected: 20877010
  full_scans: 180
```

Statement Types Per Host

```
mysql> select host,statement,total,total_latency,rows_sent from host_summary_by_statement_type
limit 5;
```

host	statement	total	total_latency	rows_sent
background	select	1	36.99 ms	1
li1306-116.members.linode.com	select	12812362	1.81 h	13354603
li1306-116.members.linode.com	update	6825577	58.96 m	0
li1306-116.members.linode.com	commit	1325481	58.02 m	0
li1306-116.members.linode.com	insert	5232533	39.13 m	0

```
5 rows in set (0.02 sec)
```

File IO

```
mysql> select * from io_global_by_file_by_bytes limit 10;
```

file	count_read	total_read	avg_read	count_write	total_written	avg_write	total	write_pct
@@datadir/xb_doublewrite	0	0 bytes	0 bytes	194812	298.55 GiB	1.57 MiB	298.55 GiB	100.00
@@datadir/mysql3-slow.log	0	0 bytes	0 bytes	88694996	67.25 GiB	814 bytes	67.25 GiB	100.00
@@datadir/tpcc5/order_line1.ibd	342876	5.23 GiB	16.00 KiB	1769489	27.00 GiB	16.00 KiB	32.23 GiB	83.77
@@datadir/tpcc1/order_line1.ibd	341504	5.21 GiB	16.00 KiB	1765234	26.94 GiB	16.00 KiB	32.15 GiB	83.79
@@datadir/tpcc3/order_line1.ibd	335372	5.12 GiB	16.00 KiB	1770525	27.02 GiB	16.00 KiB	32.13 GiB	84.07
@@datadir/tpcc4/order_line1.ibd	327280	4.99 GiB	16.00 KiB	1708848	26.07 GiB	16.00 KiB	31.07 GiB	83.93
@@datadir/tpcc2/order_line1.ibd	320811	4.90 GiB	16.00 KiB	1619925	24.72 GiB	16.00 KiB	29.61 GiB	83.47
@@datadir/tpcc5/stock1.ibd	398634	6.08 GiB	16.00 KiB	1370460	20.91 GiB	16.00 KiB	26.99 GiB	77.47
@@datadir/tpcc2/stock1.ibd	397969	6.07 GiB	16.00 KiB	1371018	20.92 GiB	16.00 KiB	26.99 GiB	77.50
@@datadir/tpcc1/stock1.ibd	398392	6.08 GiB	16.00 KiB	1368650	20.88 GiB	16.00 KiB	26.96 GiB	77.45

```
10 rows in set (0.02 sec)
```

IO by Event

```
mysql> select event_name, total_latency, total_read, total_written from io_global_by_wait_by_latency limit 10;
```

event_name	total_latency	total_read	total_written
innodb/innodb_log_file	7.27 h	11.69 MiB	42.61 GiB
innodb/innodb_data_file	1.06 h	86.82 GiB	298.79 GiB
sql/binlog	59.84 m	1.80 GiB	17.64 GiB
sql/slow_log	25.70 m	0 bytes	67.30 GiB
innodb/innodb_parallel_dblwrite_file	7.71 m	3.75 MiB	298.78 GiB
innodb/innodb_temp_file	29.52 s	592.00 KiB	667.45 MiB
sql/io_cache	1.39 s	256.50 MiB	256.50 MiB
sql/binlog_index	28.18 ms	3.95 KiB	0 bytes
sql/pid	3.02 ms	0 bytes	6 bytes
sql/ERRMSG	2.05 ms	273.98 KiB	0 bytes

```
10 rows in set (0.01 sec)
```

Memory Allocation

```
mysql> select event_name, current_count, current_alloc from memory_global_by_current_bytes limit 10;
```

event_name	current_count	current_alloc
memory/innodb/buf_buf_pool	4	533.00 MiB
memory/temptable/physical_ram	95	95.00 MiB
memory/sql/TABLE	19757	59.73 MiB
memory/performance_schema/events_statements_summary_by_digest	1	39.67 MiB
memory/performance_schema/events_errors_summary_by_thread_by_error	257	34.38 MiB
memory/innodb/ut0link_buf	2	24.00 MiB
memory/performance_schema/table_handles	2	18.12 MiB
memory/innodb/memory	10443	17.96 MiB
memory/performance_schema/events_errors_summary_by_user_by_error	129	17.19 MiB
memory/performance_schema/events_errors_summary_by_host_by_error	129	17.19 MiB

```
10 rows in set (0.01 sec)
```

Enhanced Metrics Table

```
mysql> select * from metrics limit 10;
```

Variable_name	Variable_value	Type	Enabled
aborted_clients	154	Global Status	YES
aborted_connects	446	Global Status	YES
acl_cache_items_count	0	Global Status	YES
binlog_cache_disk_use	0	Global Status	YES
binlog_cache_use	4414737	Global Status	YES
binlog_snapshot_file	mysql3-bin.000163	Global Status	YES
binlog_snapshot_position	871448732	Global Status	YES
binlog_stmt_cache_disk_use	0	Global Status	YES
binlog_stmt_cache_use	382	Global Status	YES
bytes_received	17290791216	Global Status	YES

```
10 rows in set (0.04 sec)
```

Processlist – Includes Background

```
mysql> select * from processlist limit 1 \G
***** 1. row *****
      thd_id: 30
      conn_id: NULL
      user: innodb/fts_optimize_thread
      db: NULL
      command: NULL
      state: NULL
      time: 229883
current_statement: NULL
statement_latency: NULL
  progress: NULL
  lock_latency: NULL
  rows_examined: NULL
  rows_sent: NULL
  rows_affected: NULL
  tmp_tables: NULL
  tmp_disk_tables: NULL
  full_scan: NO
  last_statement: NULL
last_statement_latency: NULL
  current_memory: 1.91 KiB
  last_wait: NULL
last_wait_latency: NULL
  source: NULL
  trx_latency: NULL
  trx_state: NULL
  trx_autocommit: NULL
  pid: NULL
  program name: NULL
1 row in set (0.30 sec)
```


Do not run out of auto_increment space

```
mysql> select * from schema_auto_increment_columns \G
***** 1. row *****
      table_schema: sbtest
      table_name: sbtest1
      column_name: id
      data_type: int
      column_type: int(11)
      is_signed: 1
      is_unsigned: 0
      max_value: 2147483647
      auto_increment: 1000000
      auto_increment_ratio: 0.0005
1 row in set (0.01 sec)
```

Redundant Indexes

```
mysql> select * from schema_redundant_indexes \G
***** 1. row *****
      table_schema: sbtest
      table_name: sbtest1
      redundant_index_name: k_1
      redundant_index_columns: k
      redundant_index_non_unique: 1
      dominant_index_name: k
      dominant_index_columns: k,c
      dominant_index_non_unique: 1
      subpart_exists: 0
      sql_drop_index: ALTER TABLE `sbtest`.`sbtest1` DROP INDEX
`k_1`
1 row in set (0.05 sec)
```

Table Access Statistics

```
mysql> select * from schema_table_statistics limit 5 \G
*****1. row*****
  table schema: tpcc1
  table name: new orders1
  total_latency: 15.37 m
  rows_fetched: 857480
  fetch_latency: 7.19 m
  rows_inserted: 288180
  insert_latency: 8.02 m
  rows_updated: 0
  update_latency: 0 ps
  rows_deleted: 284670
  delete_latency: 9.51 s
  io_read_requests: 246
    io_read: 3.89 MiB
  io_read_latency: 210.84 ms
  io_write_requests: 8041
    io_write: 125.64 MiB
  io_write_latency: 3.39 s
  io_misc_requests: 6873
  io_misc_latency: 1.68 s
```

Unused Indexes

```
mysql> select * from schema_unused_indexes where  
object_schema='tpcc1';
```

object_schema	object_name	index_name
tpcc1	history1	fkey_history_11
tpcc1	order_line1	fkey_order_line_21
tpcc1	stock1	fkey_stock_21

```
3 rows in set, 1 warning (0.02 sec)
```

Are you using SSL ?

```
mysql> select * from session_ssl_status limit 10;
```

thread_id	ssl_version	ssl_cipher	ssl_sessions_reused
12152	TLSv1.1	DHE-RSA-AES256-SHA	0
12158			0
12159			0
12160			0
12005	TLSv1.1	DHE-RSA-AES256-SHA	0
12006	TLSv1.1	DHE-RSA-AES256-SHA	0
12007	TLSv1.1	DHE-RSA-AES256-SHA	0
12008	TLSv1.1	DHE-RSA-AES256-SHA	0
12011	TLSv1.1	DHE-RSA-AES256-SHA	0
12012	TLSv1.1	DHE-RSA-AES256-SHA	0

```
10 rows in set (0.00 sec)
```

Statements info

```
mysql> select * from statement_analysis limit 5,1 \G
***** 1. row *****
      query: SHOW GLOBAL STATUS
        db: NULL
      full_scan: *
    exec_count: 240568
      err_count: 0
     warn_count: 0
  total_latency: 20.88 m
    max_latency: 153.12 ms
    avg_latency: 5.21 ms
   lock_latency: 1.70 m
      rows_sent: 117878320
  rows_sent_avg: 490
 rows_examined: 235756640
rows_examined_avg: 980
   rows_affected: 0
rows_affected_avg: 0
      tmp_tables: 240568
   tmp_disk_tables: 0
      rows_sorted: 0
sort_merge_passes: 0
      digest: 070e38632eb4444e50cdcbf0b17474ba801e203add89783a24584951442a2317
    first_seen: 2019-06-11 19:01:32.245462
    last_seen: 2019-06-14 13:50:07.158611
1 row in set (0.01 sec)
```

Errors and Warnings ?

```
mysql> select * from statements_with_errors_or_warnings limit 1 \G
***** 1. row *****
      query: SELECT `no_o_id` FROM `new_ord ... o_o_id` ASC LIMIT ? FOR UPDATE
         db: tpcc1
    exec_count: 298942
        errors: 7
    error_pct: 0.0023
    warnings: 0
warning_pct: 0.0000
  first_seen: 2019-06-11 19:02:31.104121
  last_seen: 2019-06-14 13:52:57.834563
    digest: 366d8c4aa612adf0ec3e06cb020fd395e1da1fd4f6b676530c214771aed00b55
1 row in set (0.01 sec)
```

User Info

```
mysql> select * from user_summary limit 1 \G
***** 1. row *****
      user: appl
    statements: 195012171
 statement_latency: 1.38 d
statement_avg_latency: 612.32 us
      table_scans: 1428
        file_ios: 226986594
   file_io_latency: 4.96 h
current_connections: 36
  total_connections: 476
    unique_hosts: 3
   current_memory: 125.39 MiB
total_memory_allocated: 9.13 TiB
1 row in set (0.08 sec)
```


More Wait Info (If Enabled)

```
mysql> select  events, total, total_latency  from waits_global_by_latency limit 10;
```

events	total	total_latency
wait/io/file/innodb/innodb_log_file	108966439	9.17 h
wait/io/table/sql/handler	2615937816	5.41 h
wait/io/file/innodb/innodb_data_file	34980007	1.35 h
wait/io/file/sql/binlog	11738591	1.25 h
wait/io/file/sql/slow_log	109842706	32.11 m
wait/io/file/innodb/innodb_parallel_dblwrite_file	483426	9.70 m
wait/lock/table/sql/handler	101720352	2.82 m
wait/io/file/innodb/innodb_temp_file	24212	39.62 s
wait/io/file/sql/io_cache	18961	2.29 s
wait/io/file/sql/binlog_index	731	38.33 ms

```
10 rows in set (0.01 sec)
```

Latency Distributions

```
mysql> select * from
events_statements_histogram_global limit 1 \G
*****
1. row
*****

      BUCKET_NUMBER: 0
      BUCKET_TIMER_LOW: 0
      BUCKET_TIMER_HIGH: 10000000
      COUNT_BUCKET: 162125
COUNT_BUCKET_AND_LOWER: 162125
      BUCKET_QUANTILE: 0.001473
1 row in set (0.00 sec)
```

Logs

Error Log

General Query Log

Slow Query Log

Percona Server supports Sampling

Slow Query Log (Percona Server)

```
# Time: 2019-06-14T14:07:29.062441Z
# User@Host: app1[app1] @ li534-125.members.linode.com [50.116.47.125] Id: 7442
# Schema: tpcc1 Last_errno: 0 Killed: 0
# Query_time: 0.000635 Lock_time: 0.000123 Rows_sent: 0 Rows_examined: 0 Rows_affected: 1
# Bytes_sent: 11 Tmp_tables: 0 Tmp_disk_tables: 0 Tmp_table_sizes: 0
# InnoDB_trx_id: 366322C
# Full_scan: No Full_join: No Tmp_table: No Tmp_table_on_disk: No
# Filesort: No Filesort_on_disk: No Merge_passes: 0
# InnoDB_IO_r_ops: 0 InnoDB_IO_r_bytes: 0 InnoDB_IO_r_wait: 0.000000
# InnoDB_rec_lock_wait: 0.000000 InnoDB_queue_wait: 0.000000
# InnoDB_pages_distinct: 7
SET timestamp=1560521249;
INSERT INTO order_line1
      (ol_o_id, ol_d_id, ol_w_id, ol_number, ol_i_id, ol_supply_w_id,
ol_quantity, ol_amount, ol_dist_info)
      VALUES (181228,6,1,9,31509,1,3,169,'xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx');;
```

MySQL 8.0.14+ Extends Slow Query Log

```
# Time: 2019-06-14T14:14:22.980797Z
# User@Host: root[root] @ localhost [] Id:      8
# Query time: 0.005342  Lock time: 0.000451 Rows sent: 33
Rows examined: 197 Thread id: 8 Errno: 0 Killed: 0
Bytes received: 0 Bytes sent: 664 Read first: 1 Read last:
0 Read key: 71 Read_next: 127 Read_prev: 0 Read_rnd: 33
Read_rnd_next:
34 Sort merge passes: 0 Sort range count: 0 Sort_rows: 33
Sort scan count: 1 Created tmp disk tables: 0
Created tmp tables: 1 Start: 2019-06-14T14:14:22.975455Z
End: 2019-06-14T14:14:22.980797Z
SET timestamp=1560521662;
show tables;
```

EXPLAIN

Understand Query Execution Plan

Essential Skill for Developers and DBAs

Multiple Output Formats

Can get EXPLAIN plan for actual running query

Basic Explain

```
mysql> explain SELECT o_c_id          FROM orders1      WHERE o_id = 180962      AND o_d_id = 10
AND o_w_id = 1 \G
***** 1. row *****
      id: 1
  select_type: SIMPLE
        table: orders1
    partitions: NULL
         type: const
possible_keys: PRIMARY,idx_orders1
         key: PRIMARY
      key_len: 7
         ref: const,const,const
        rows: 1
   filtered: 100.00
      Extra: NULL
1 row in set, 1 warning (0.00 sec)
```

JSON Explain

```
mysql> explain format=json SELECT o_c_id
      WHERE o_id = 180962
      FROM orders1
      AND o_w_id = 1 \G
      AND o_d_id = 10
***** 1. row *****
EXPLAIN: {
  "query_block": {
    "select_id": 1,
    "cost_info": {
      "query_cost": "1.00"
    },
  },
  "table": {
    "table_name": "orders1",
    "access_type": "const",
    "possible_keys": [
      "PRIMARY",
      "idx_orders1"
    ],
  },
  "key": "PRIMARY",
  "used_key_parts": [
    "o_w_id",
    "o_d_id",
    "o_id"
  ],
  "key_length": "7",
  "ref": [
    "const",
    "const",
    "const"
  ],

```

```
    ],
    "rows_examined_per_scan": 1,
    "rows_produced_per_join": 1,
    "filtered": "100.00",
    "cost_info": {
      "read_cost": "0.00",
      "eval_cost": "0.10",
      "prefix_cost": "0.00",
      "data_read_per_join": "24"
    },
    "used_columns": [
      "o_id",
      "o_d_id",
      "o_w_id",
      "o_c_id"
    ]
  },
}
}
1 row in set, 1 warning (0.02 sec)
```


TREE Format (8.0.16+)

EXPLAIN: -> Aggregate: count(0)

-> Filter: ((t1.c1 + t1.c_ytd_payment) <> t1.sm)

-> Table scan on t1

-> Materialize

-> Index lookup on c using PRIMARY (c_w_id=1)

-> Select #3 (subquery in projection; dependent)

-> Aggregate: sum(order_line1.ol_amount)

-> Nested loop inner join

-> Index lookup on orders1 using idx_orders1 (o_w_id =c.c_w_id, o_d_id=c.c_d_id, o_c_id=c.c_id)

-> Filter: ((order_line1.ol_d_id = orders1.o_d_id) and (order_line1.ol_w_id = orders1.o_w_id) and
(order_line1.ol_delivery_d is not null))

-> Index lookup on order_line1 using PRIMARY (ol_w_id=c.c_w_id, ol_d_id=c.c_d_id,
ol_o_id=orders1.o_id)

EXPLAIN ANALYZE

```
mysql> EXPLAIN ANALYZE SELECT * FROM t1 JOIN t2 ON (t1.c1 = t2.c2)\G
***** 1. row *****
EXPLAIN: -> Inner hash join (t2.c2 = t1.c1) (cost=4.70 rows=6)
(actual time=0.032..0.035 rows=6 loops=1)
    -> Table scan on t2 (cost=0.06 rows=6)
(actual time=0.003..0.005 rows=6 loops=1)
    -> Hash
        -> Table scan on t1 (cost=0.85 rows=6)
(actual time=0.018..0.022 rows=6 loops=1)
```

<https://www.percona.com/blog/2019/10/28/using-explain-analyze-in-mysql-8/>

EXPLAIN FOR CONNECTION

```
mysql [localhost] {msandbox} ((none)) > explain for connection 9 (
***** 1. row *****
      id: 1
    select_type: SIMPLE
        table: employees
    partitions: NULL
         type: ALL
possible_keys: NULL
         key: NULL
        key_len: NULL
         ref: NULL
         rows: 299540
    filtered: 100.00
       Extra: NULL
***** 2. row *****
      id: 1
    select_type: SIMPLE
        table: salaries
    partitions: NULL
         type: ALL
possible_keys: NULL
         key: NULL
        key_len: NULL
         ref: NULL
         rows: 2803840
    filtered: 100.00
       Extra: Using where; Using join buffer (Block Nested Loop)
```

Tracing Optimizer

```
# Turn tracing on (it's off by default):  
SET optimizer_trace="enabled=on";  
SELECT ...;  
# your query here  
SELECT * FROM  
INFORMATION_SCHEMA.OPTIMIZER_TRACE;  
# possibly more queries...  
# When done with tracing, disable it:  
SET optimizer_trace="enabled=off";
```

<https://dev.mysql.com/doc/internals/en/optimizer-tracing.html>

Percona Monitoring and Management

Open Source Database Focused
Observability Solution from Percona

100% Free and Open Source

<http://per.co.na/PMM>



Demo Time

DEMO

Want to learn more about PMM



Come to Percona Booth to see the Demo and get your questions answered



Our team can also can help you to install PMM



If you have PMM installed already we can give you some tips on how to use it to improve your database

Custom Dashboards shown in the Demo

<https://www.percona.com/blog/2020/11/02/understanding-mysql-memory-usage-with-performance-schema/>

<https://www.percona.com/blog/2021/04/22/understanding-processes-running-on-linux-host-with-percona-monitoring-and-management/>

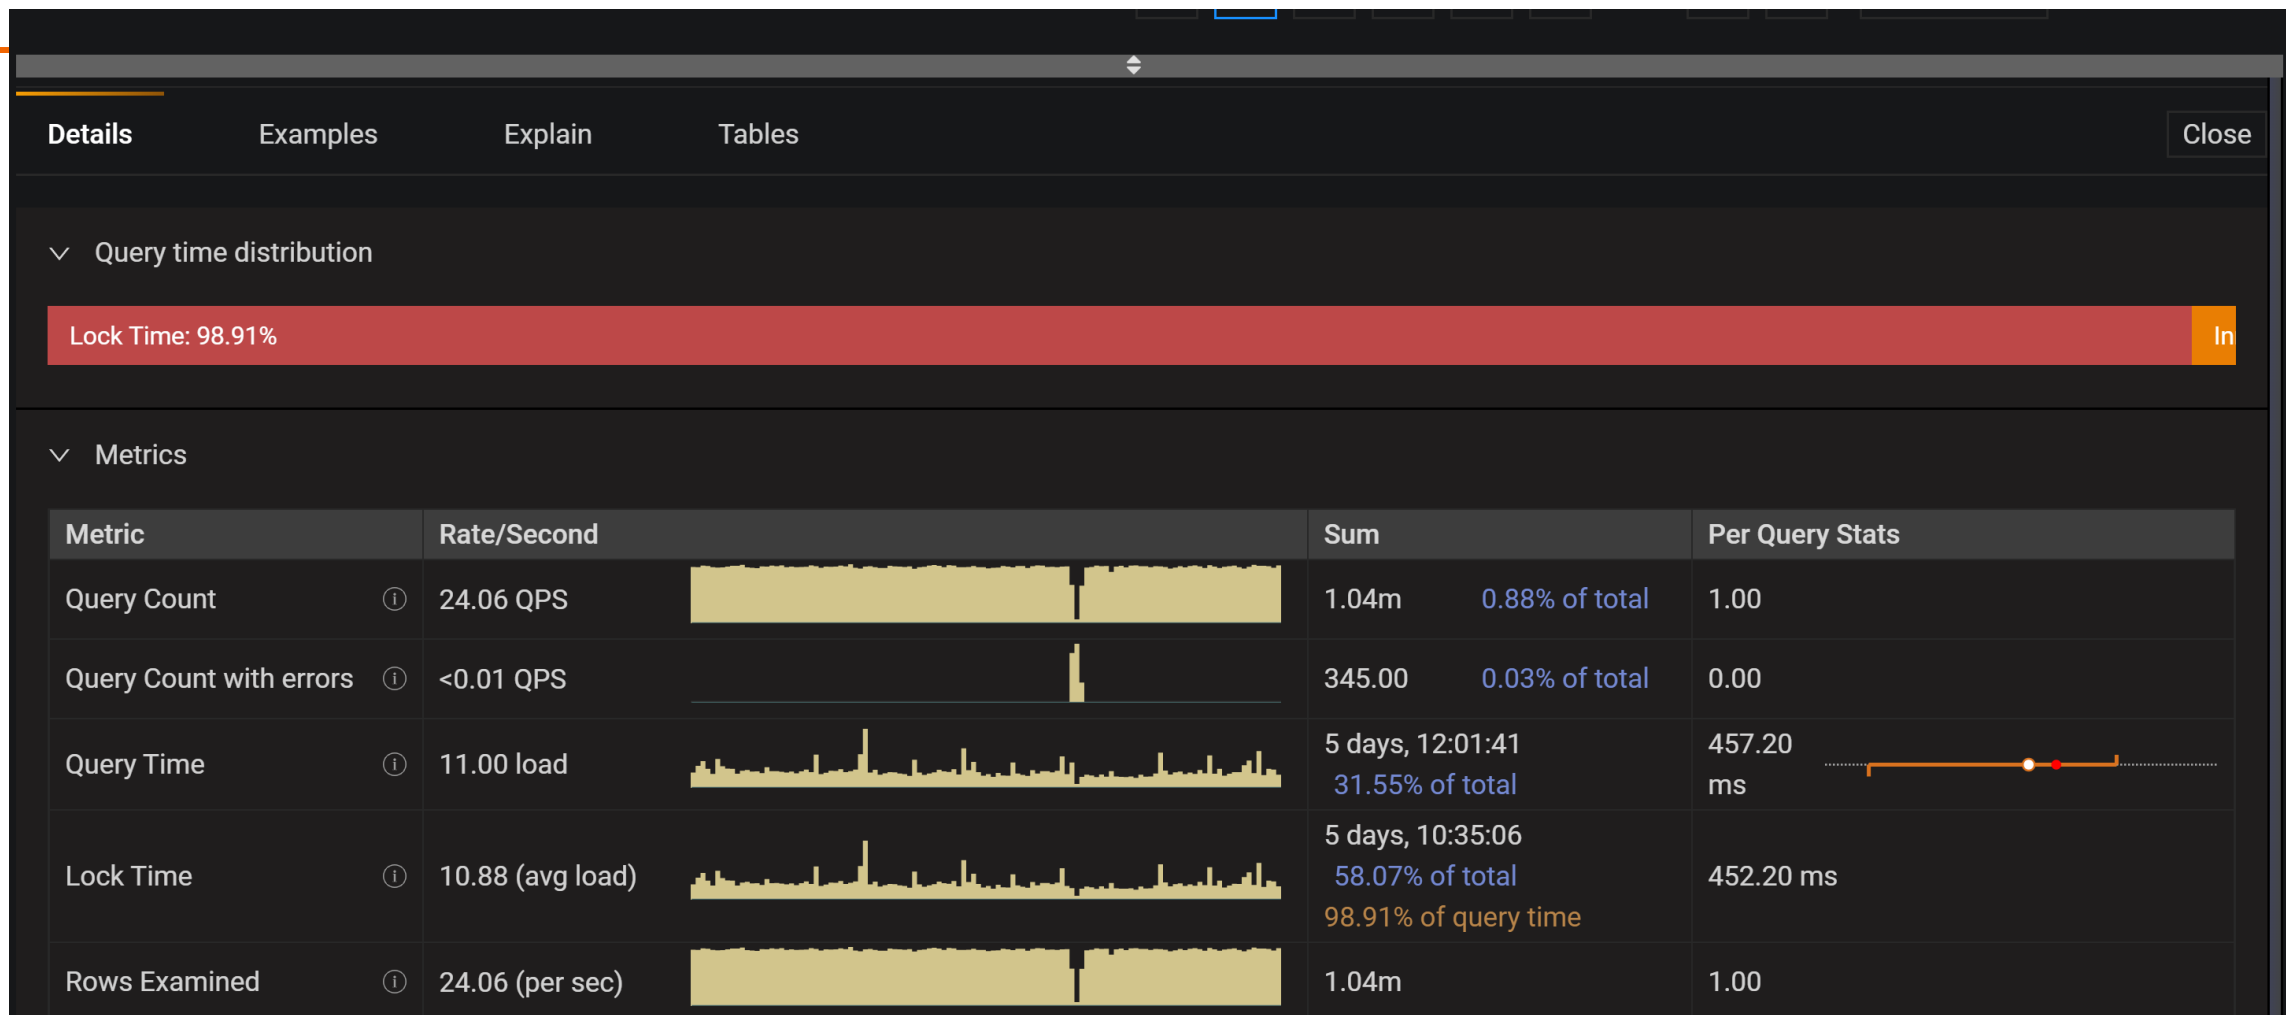
<https://www.percona.com/blog/2020/07/15/mysql-query-performance-troubleshooting-resource-based-approach/>

<https://www.percona.com/blog/2020/06/17/red-method-for-mysql-performance-analyses/>

Which Queries are Causing the Load



Why Are they Causing this Load



How to Improve their Performance

▼Example

```
SELECT DISTINCT c
FROM sbtest1
WHERE id
      BETWEEN 5559
      AND 5658
ORDER BY c
```

▼CREATE

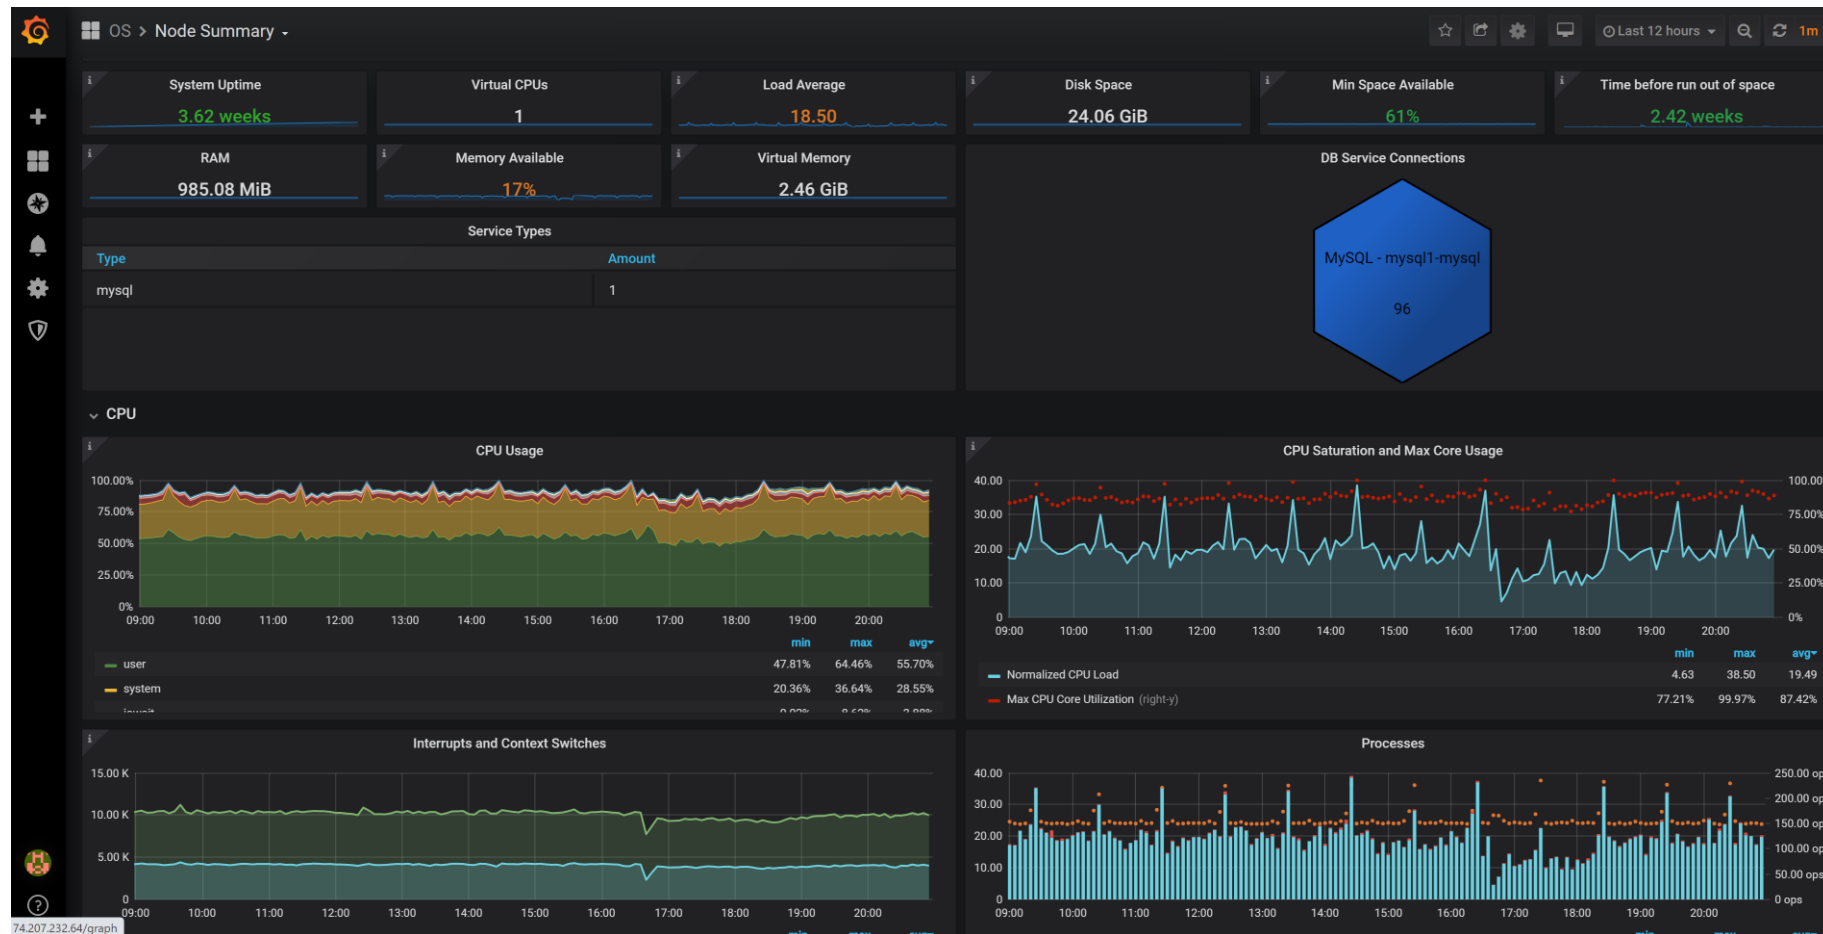
```
CREATE TABLE `sbtest1` (
  `id` int(10) unsigned NOT NULL AUTO_INCREMENT,
  `k` int(10) unsigned NOT NULL DEFAULT '0',
  `c` char(120) NOT NULL DEFAULT '',
  `pad` char(60) NOT NULL DEFAULT '',
  PRIMARY KEY (`id`),
  KEY `k_1` (`k`)
) ENGINE=InnoDB AUTO_INCREMENT=100000001 DEFAULT
```

▼JSON

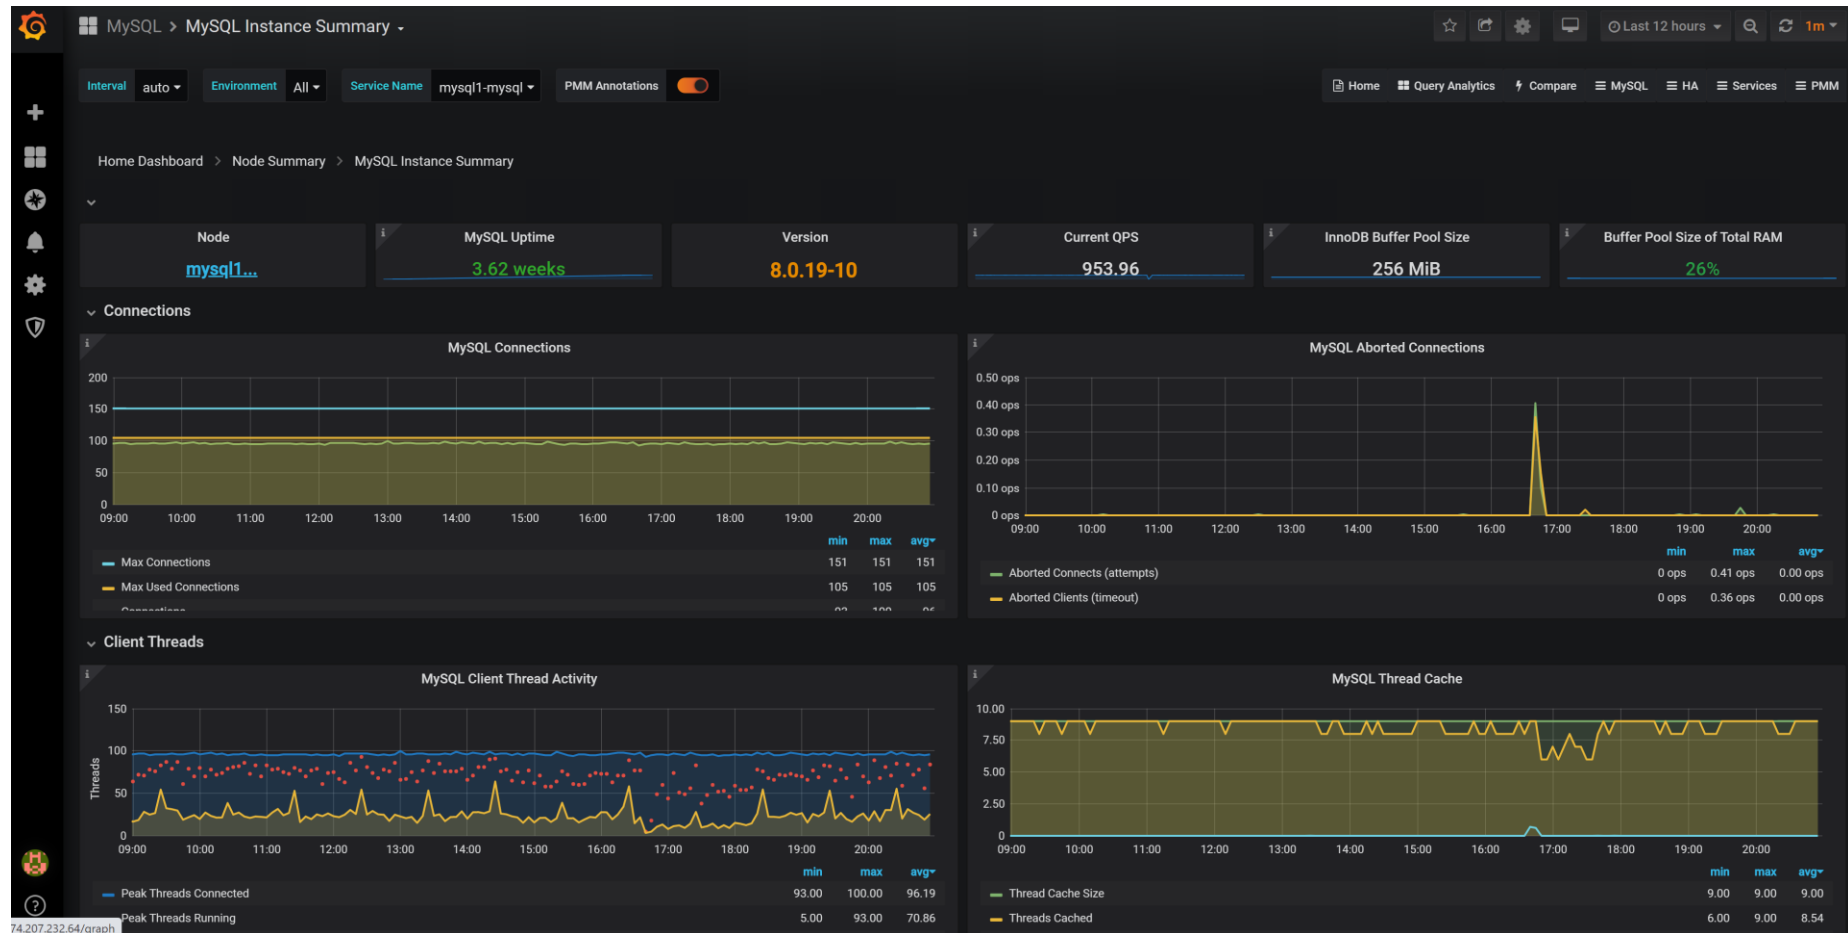
Expand All

```
-{
  "query_block": -{
    "select_id": 1,
    "cost_info": +{...},
    "ordering_operation": -{
      "using_filesort": false,
      "duplicates_removal": -{
        "using_temporary_table": true,
        "using_filesort": true,
        "cost_info": +{...},
        "table": +{...}
      }
    }
  }
}
```

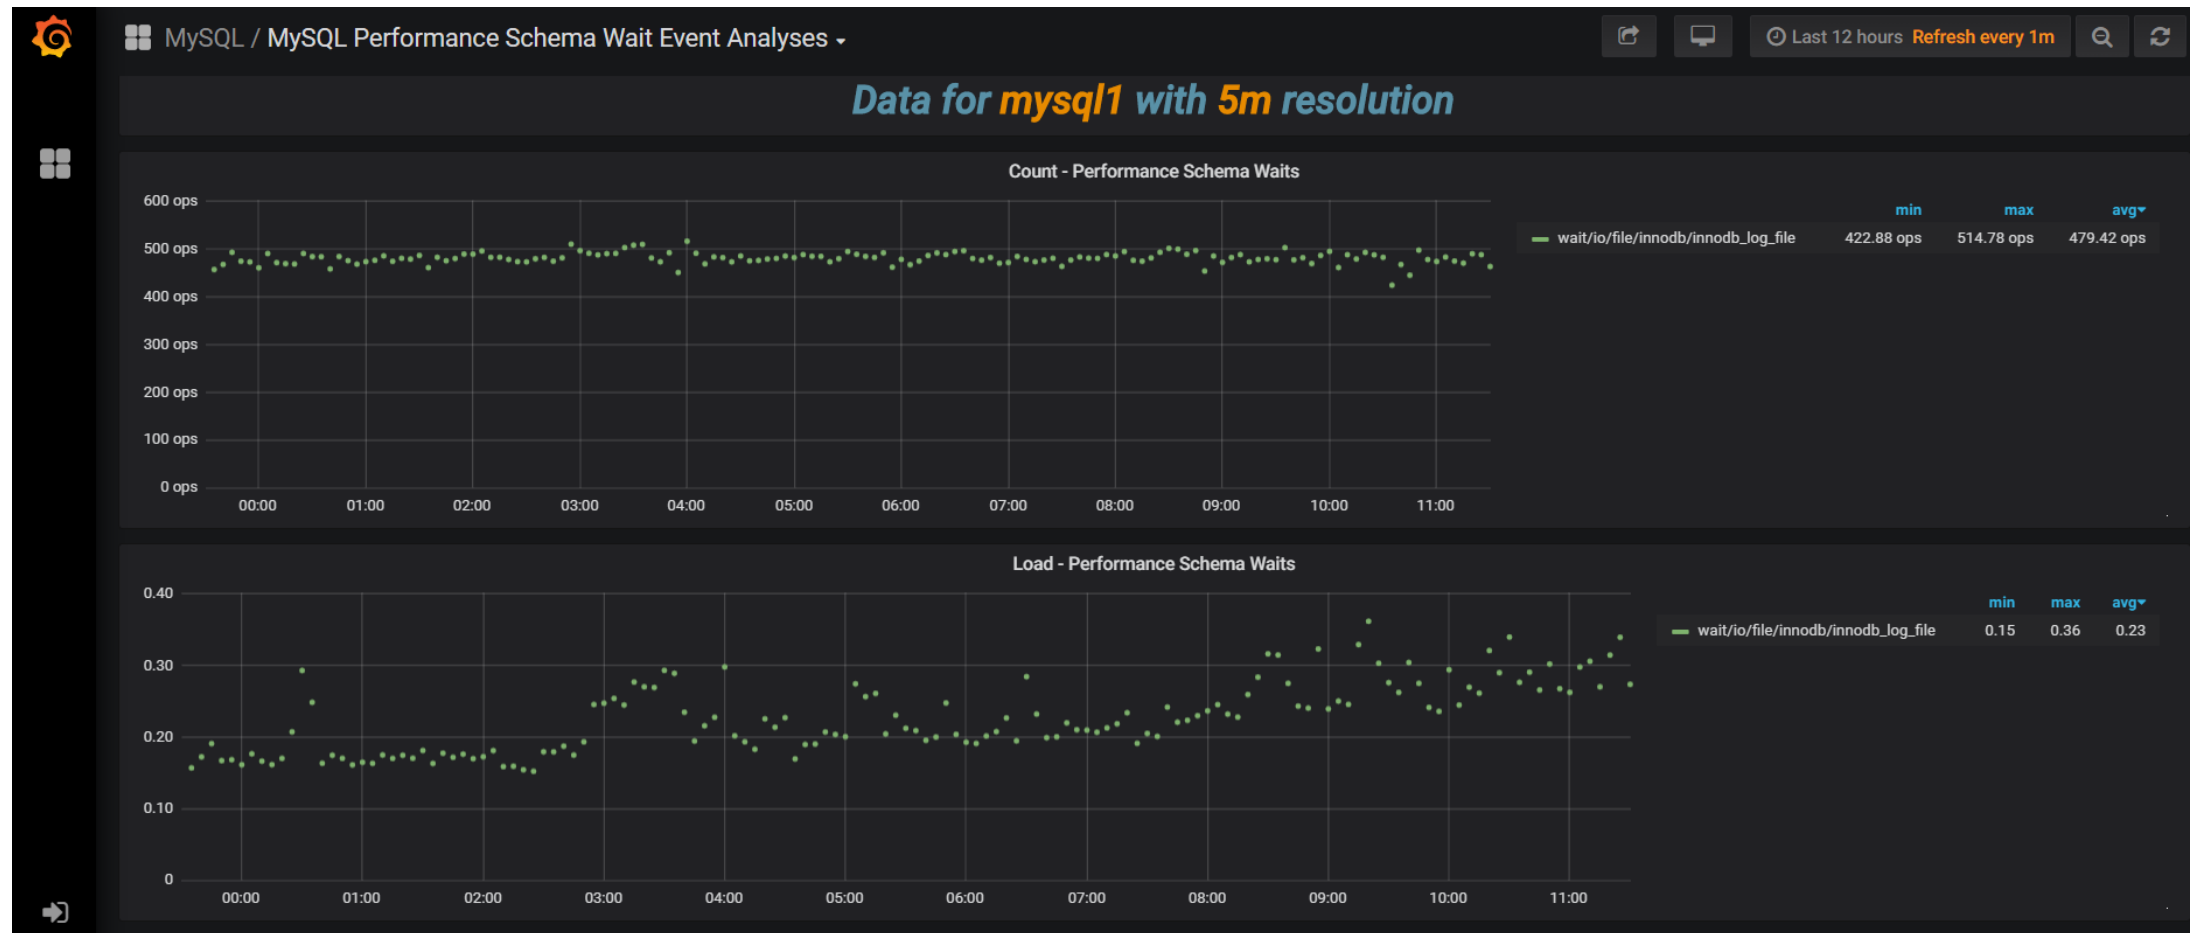
System Monitoring



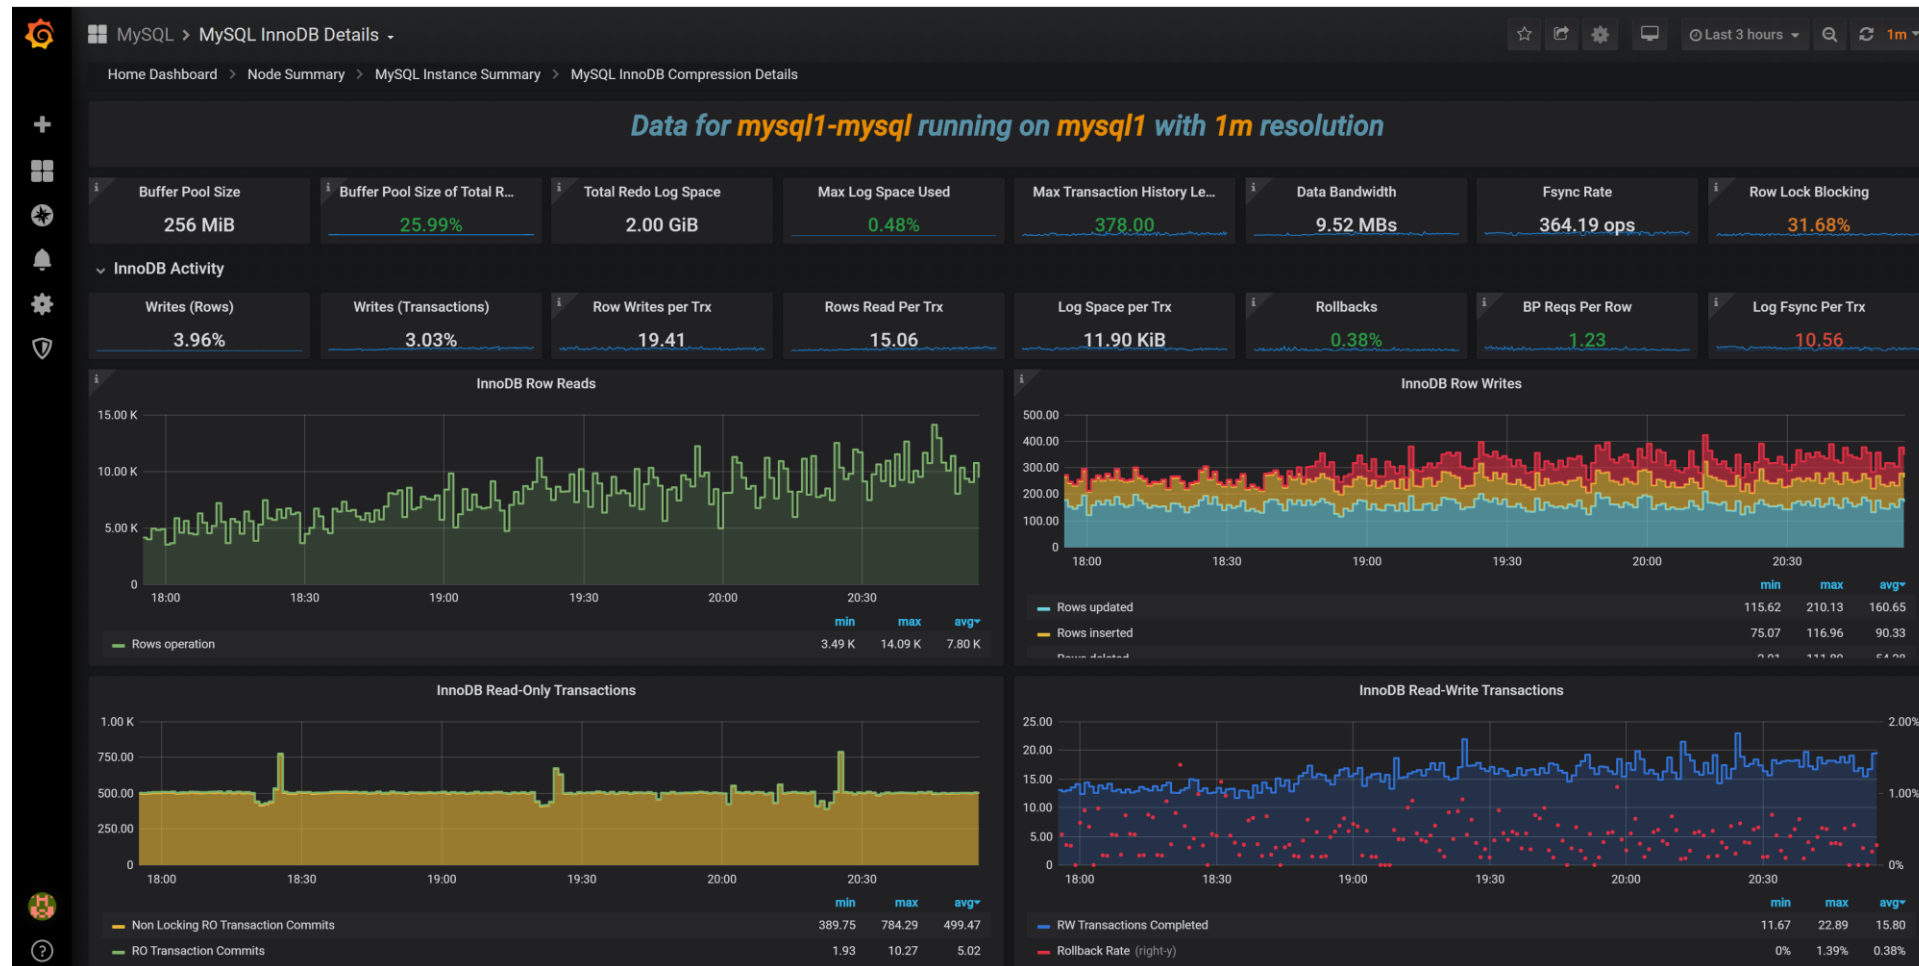
MySQL Instance Summary



Advanced Performance Schema Analyses



InnoDB Insights



See it Live!

<https://pmmdemo.percona.com>

Thank You!

Twitter: @percona @peterzaitsev
